4. POTENTIAL DEVICES AND METHODS USEFUL IN RADIOS EMPLOYING RF AND IF DIGITIZATION

This section describes some methods and devices that may be useful in implementing radio receivers where digitization occurs at the RF or IF. These methods and devices probably will be considered in the design of future radio receivers as the digitization of received signals progresses toward the receive antenna. Various quantization techniques are discussed in detail since the authors believe that this area has great potential for use in radio receiver applications. Nonlinear compression devices are discussed, including log amplifiers and automatic gain control devices. Postdigitization algorithms for improving the SFDR provide extended dynamic range capability for presently available ADC's. Sampling downconverter technology is introduced. This technology is based on the theory of bandpass sampling and may be coupled with ADC technology to provide improved ADC performance for bandpass sampling applications. Finally, specialized integrated circuits for digital signal-processing tasks required in radio receiver applications are mentioned. An example of this type of technology, the Harris HSP50016 digital downconverter, is presented.

4.1 Quantization Techniques

Section 4.1, Quantization Techniques (from pages 36-47), is currently unavailable for distribution over the Internet. Please contact Jeff Wepman via e-mail at jwepman@its.bldrdoc.gov for further information.

Section 4.1, Quantization Techniques (from pages 36-47), is currently unavailable for distribution over the Internet. Please contact Jeff Wepman via e-mail at jwepman@its.bldrdoc.gov for further information.

Section 4.1, Quantization Techniques (from pages 36-47), is currently unavailable for distribution over the Internet. Please contact Jeff Wepman via e-mail at jwepman@its.bldrdoc.gov for further information.

Section 4.1, Quantization Techniques (from pages 36-47), is currently unavailable for distribution over the Internet. Please contact Jeff Wepman via e-mail at jwepman@its.bldrdoc.gov for further information.

Section 4.1, Quantization Techniques (from pages 36-47), is currently unavailable for distribution over the Internet. Please contact Jeff Wepman via e-mail at jwepman@its.bldrdoc.gov for further information.

Section 4.1, Quantization Techniques (from pages 36-47), is currently unavailable for distribution over the Internet. Please contact Jeff Wepman via e-mail at jwepman@its.bldrdoc.gov for further information.

Section 4.1, Quantization Techniques (from pages 36-47), is currently unavailable for distribution over the Internet. Please contact Jeff Wepman via e-mail at jwepman@its.bldrdoc.gov for further information.

Section 4.1, Quantization Techniques (from pages 36-47), is currently unavailable for distribution over the Internet. Please contact Jeff Wepman via e-mail at jwepman@its.bldrdoc.gov for further information.

Section 4.1, Quantization Techniques (from pages 36-47), is currently unavailable for distribution over the Internet. Please contact Jeff Wepman via e-mail at jwepman@its.bldrdoc.gov for further information.

Section 4.1, Quantization Techniques (from pages 36-47), is currently unavailable for distribution over the Internet. Please contact Jeff Wepman via e-mail at jwepman@its.bldrdoc.gov for further information.

Section 4.1, Quantization Techniques (from pages 36-47), is currently unavailable for distribution over the Internet. Please contact Jeff Wepman via e-mail at jwepman@its.bldrdoc.gov for further information.

4.2 Nonlinear Devices for Amplitude Compression

As discussed in Section 4.1.2, logarithmic compression at the input to the quantizer can improve the SNR by making it independent of the input signal variance. In actual implementation, where system noise defines a minimum signal input amplitude, this form of compression can be achieved through the use of log amplifiers. Likewise, by using variable gain, it is possible to adapt the signal amplitude to the characteristics of the quantizer. This, too, can be implemented by using automatic gain control (AGC) amplifiers. Both log amplifiers and AGC amplifiers are discussed in greater detail below [30-33].

4.2.1 Log Amplifiers

There are several different types of log amplifiers, each suited to different applications. Common to all is some form of logarithmic compression of signal parameters. A simple operational amplifier circuit that uses the nonlinear (logarithmic) characteristics of a p-n junction can be used [34]; however, these circuits suffer from temperature variations, limited dynamic range, and slow rise times. Depending on the application, these circuits may or may not be adequate. A more powerful technique is the approximation of the logarithmic function using the summation of linear (or curved) lines. Typically these are implemented with differential amplifiers using integrated circuit technology and may be purchased as discrete components or built into the front-end receiver circuitry. There are three basic types of these log amplifiers: detector log video amplifiers, successive detection log amplifiers, and true log IF amplifiers.

Detector log video amplifiers are suited to applications where phase and frequency information is not necessary. The envelope of the input signal simply is converted to a log-compressed video signal at the output. Total input dynamic range of these amplifiers is generally 50 dB. A typical application might be the demodulation and logarithmic compression of an amplitude-modulated signal.

The successive detection log amplifier provides two outputs. One is the same as the detector log video amplifier output, described above. The other is a limited IF signal. The former provides amplitude information expressed in logarithmic form and the latter provides phase and frequency information. The limited IF signal is a copy of the input signal except that its amplitude variation is compressed and limited by a transfer function similar to that in Figure 23. This figure shows the limited IF output power as a function of the input signal power. Typically these amplifiers have a total input dynamic range of 80 dB. A successive detection log amplifier recently has been announced by Microphase, Inc. that provides a 100-dB input dynamic range.

The "true" logarithmic IF amplifier is called so because the IF output is a bipolar logarithmic function of the IF input signal (without any limiting). This amplifier may also have an output that is the same as the detector log video amplifier. Due to the dual-polarity of the output signal, these amplifiers function well in logarithmic IF applications. As might be expected, there is some deviation from the true logarithmic curve (toward a more linear function) as signals approach zero, but this is not generally a problem if the system noise power is set to the minimum signal input amplitude of the log amplifier (typically -80 dBm). These amplifiers also generally have a total input dynamic range of 80 dB. In addition, they inherently have low phase shifts over wide variations in input signal power.

Logarithmic amplifiers have wide instantaneous dynamic range. Instantaneous dynamic range means that at any point in time, all signals within the dynamic range of the amplifier appear at the

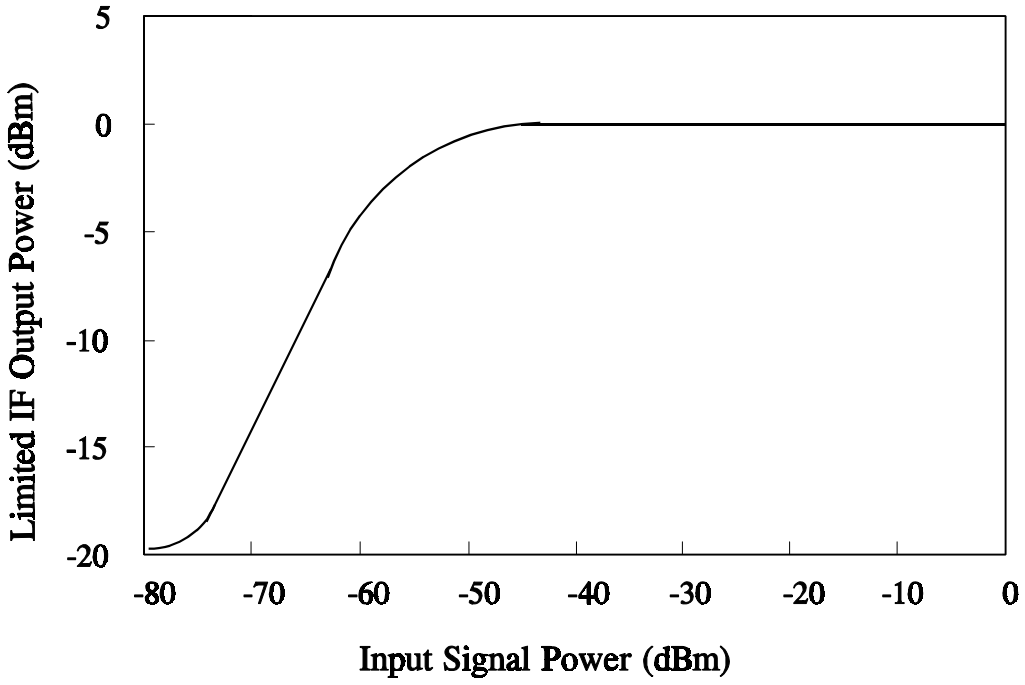

Figure 23. Limited IF output power as a function of the input signal power.

output. Take for an example, an input containing two signals, one at -5 dBm and one at -75 dBm. For a logarithmic amplifier with an 80-dB dynamic range, it is possible to maintain the integrity of both signals at the output, but in a compressed form. Instead of the log amplifier, consider a variable attenuator placed before a fixed-gain amplifier. With the same input consisting of both the -5-dBm and -75-dBm signals, assume that the fixed-gain amplifier would saturate if no attenuation was used. Therefore, the attenuation of the variable attenuator must be increased to prevent amplifier saturation. In doing so, the noise figure of the combination of the variable attenuator and the fixed-gain amplifier will increase. This may cause the smaller signal (-75 dBm) to be overpowered by the noise and go undetected. Logarithmic amplifiers are, therefore, good for processing multiple signals where small signals are present simultaneously with large signals. Logarithmic amplifiers must be used carefully in receiver systems, however. Being nonlinear devices, they may cause distortion of the input signals. A careful analysis of the effects of distortion on the desired received signals must be performed.

4.2.2 Automatic Gain Control

AGC is used to normalize the output signal level despite variations in input signal power. A typical AGC device operates using a feedback loop whereby the power of the output of a variable gain device (amplifier or attenuator) is sampled and used to adjust the gain of the variable gain device itself. In this way, the output power is maintained at a relatively constant level. The response time of the AGC loop is perhaps the most critical design parameter and is dependent upon the application. Generally the AGC device is adjusted to react fast enough to normalize the overall power but slow enough to maintain the desired information content of the signal. Depending upon the application, the control voltage for the variable gain device may be used to decompress the output signal after digitization.

Typically AGC devices can have a total dynamic range of 80 dB. But because they operate by adjusting the gain of the system, the instantaneous dynamic range is much lower than that for logarithmic amplifiers. Therefore, these systems generally are better suited for narrower bandwidth systems where individual signals are isolated by appropriate filtering before the AGC. AGC's, unlike logarithmic amplifiers, are not effective in receivers where weak desired signals need to be detected in the presence of strong undesired signals.

AGC devices can be purchased as discrete components or built into the front-end receiver circuitry. As with any feedback system, stability is an important consideration. In radio receiver applications that use an AGC device before an ADC, there should be enough gain in the system so that the system noise at the output of the AGC device prior to digitization is slightly greater than the quantization noise power of the digitizer. To maximize the SNR, the normalized output of the AGC device should be set as large as possible without causing the ADC to overload or generate excessive spurious responses. A more detailed discussion of AGC devices can be found in [35].

4.3 Postdigitization Algorithms for Improving Spurious Free Dynamic Range

As previously mentioned, there are numerous postdigitization techniques for optimizing the quantization process. Many of these techniques are used to increase the SNR of the quantizer by improving the predictor characteristics of differential quantization schemes. Another area of postdigitization processing provides compensation for the nonlinearities that occur in practical implementations of ADC's. As discussed in Section 2 of this report, these nonlinearities produce spurious signals that can reduce significantly the SFDR of the ADC. The purpose of this compensation is to suppress the spurious responses below the noise in the frequency band from 0 to $f_s/2$. Two of these techniques, phase-plane and state variable compensation, are discussed below [13,14].

Both techniques are used to identify a set of correction factors that can be used to compensate for any nonlinearity throughout the full amplitude range of the ADC. In phase-plane compensation, the procedure for correcting the digitized signal is as follows: The input signal is split into two separate signals. One signal is fed into the ADC and the other is sent through an analog differentiator and then digitized by a second ADC. The differentiated signal is used to determine the instantaneous slope of the signal. The output of both ADC's then is used to determine the correction factor to be applied to the ADC output representing the digitized input signal. A table consisting of correction factors for each possible combination of quantization level and instantaneous slope is developed for an individual ADC based on measurements of that particular ADC. This table then is stored in RAM and is used to provide the correct ADC output for any given input signal amplitude. Studies using this technique show as much as a 15- to 16-dB (about 2.5 bits) improvement in the SFDR over uncompensated ADC's [14]. This improvement, however, is restricted to a narrow frequency band well below $f_s/2$.

In an effort to improve the SFDR for all frequencies in the 0 to $f_s/2$ frequency band, a state variable compensation technique was also proposed. This type of compensation is implemented by applying the input signal to an ADC and splitting the output of the ADC into two signals. One of these signals is used without modification while the other is delayed by a single clock cycle (one sample of the input signal). The two outputs, representing the quantization levels for the present and previous ADC outputs, then are used to determine the correction factor to be applied to the present ADC output. A table of correction factors for each possible combination of present and previous quantization levels is developed for an individual ADC based on measurements of that particular ADC. As in phase-plane compensation, this table is stored in RAM and is used to provide the correct ADC output for any given input signal amplitude. Tests using this technique also show as much as a 16-dB improvement in the SFDR over the entire 0 to $f_s/2$ frequency band for the particular sampling rate.

While compensation techniques require additional hardware and testing of individual ADC's, they can improve the SFDR of the ADC significantly without increasing its resolution (number of bits). In essence, they bring the characteristics of the ADC closer to the theoretical expectation of its performance. An underlying assumption in these techniques, however, is that the ADC characteristics are static. Testing of ADC's has shown that for most ADC's this is a valid assumption [14].

4.4 Sampling Downconverters

The 6300 series sampling downconverters manufactured by Watkins-Johnson are an example of a technology potentially useful in the RF front-end of radio receivers. These sampling downconverters are microwave devices that use bandpass sampling techniques to downconvert an RF signal (typically in the 2- to 18-GHz range) to an IF signal (typically at 70 MHz up to 500 MHz). In its current configuration, the sampling downconverter functionally performs in the same manner as the conventional mixer. However, the device is actually a sample-and-hold circuit that uses a step recovery diode (SRD) to generate a sampling pulse train from a sinusoidal "clock" frequency. The sample-and-hold circuits in these downconverters are designed to work in the 2- to 18-GHz range, far higher in frequency than the sample-and-hold circuits integral to most ADC's. In the future, this technology may be coupled with high-resolution quantizers to produce bandpass sampling ADC's with a high analog input frequency capability and the high resolution of quantizers that would operate at much lower frequencies. Use of these sampling downconverters for radio receiver applications requires some careful considerations, however. The frequency content of the input must be bandlimited properly to prevent spectrum overlap in the desired signal output as with any bandpass sampling scheme. These downconverters do suffer a higher conversion loss than conventional mixers. The maximum specified conversion loss for the 6300 series sampling downconverters can be anywhere from 13-25 dB depending on the specific model. Typical conversion loss for conventional mixers ranges from roughly 5-9 dB. As with any downconverter, spurious suppression must be considered carefully. The specifications for the Watkins-Johnson 6300 series sampling downconverters state a minimum 15-dBc spurious suppression at a -10-dBm RF input level [36]. (This spurious suppression does not include the local oscillator (LO) signal leakage at the IF port or the 2^{nd} harmonic of the IF signal.) More typically, even a bad spur would be 25 dBc with the -10-dBm RF input level. Operation at lower RF input levels will provide even better spurious suppression [37]. These types of devices may be very useful for future radio receiver front-ends and ADC's.

It is tempting to compare the spurious suppression of the sampling downconverter to that of conventional mixers. While this comparison could be made for a specific mixer, it is not possible to make any general conclusions. While there are several types of conventional mixers, double-balanced mixers are the industry standard [38]. Even within the double-balanced mixer type, there are several different classes of mixers. These classes of mixers include class 1, class 2, and class 3 mixers. Class 2 mixers require more LO power and have better spurious suppression than class 1 mixers. Similarly, class 3 mixers require more LO power and have better spurious suppression than class 2 mixers. The higher the mixer class, the more LO power is required but the better the spurious suppression is. In addition to the different classes of mixers, spurious suppression is a function of LO and RF input power levels and frequencies. While there are specifications on mixers that help predict spurious suppression, spurious suppression for a specific mixer should be determined by measurement over a well-defined set of conditions [38].

4.5 Specialized Integrated Circuits

As digitization in radio receivers moves closer to the receive antenna, more of the traditionally analog receiver functions will be replaced with digital signal-processing algorithms. In addition, the area of digital communications requires extensive digital signal processing. While the speed of general digital signal processors continues to increase, specialized integrated circuits (IC's) can, in general, digitally implement radio receiver functions faster. Use of these specialized IC's should be quite prevalent in radio receiver applications of the future.

A current example of this type of specialized IC is the single chip digital downconverter. Digital downconversion offers a clear advantage over conventional analog downconversion by providing more precise frequency control of the LO and by providing a more ideal mixing operation. Current digital downconverters can control LO frequencies to within less than .01 Hz out of several MHz and can provide mixing with over 100 dB of SFDR. In addition, on-chip filters can be programmed to produce almost any desired bandwidth while preserving linear phase. In contrast to analog equivalents, these chips provide reproducible component accuracies that do not degrade over time or temperature.

The Harris HSP50016 is a programmable monolithic digital downconverter capable of extracting a narrow baseband signal from a wideband input bandpass signal (RF or IF) [39,40]. The digital downconverter multiplies a digitized wideband input signal by samples representing both an in-phase LO and a quadrature-phase LO to produce in-phase and quadrature-phase baseband signals. Next, the digital downconverter filters the in-phase and quadrature-phase signals to extract the channel of interest. It then decimates the output to a sampling rate of twice the maximum frequency of the signal of interest. The digital downconverter can also be configured to provide a downconverted version of the input bandpass signal at a new IF. The HSP50016 can operate on 16-bit data at a rate of up to 52 Msamples/s [41].